

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.


**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**



#2

INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed


Dated

10 JUL 2001

BEST AVAILABLE COPY

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)



The Patent Office

Cardiff Road
Newport

Gwent NP9 1RH

1. Your reference

GW-G30244

2. Patent application number
(The Patent Office will assign a number if you do not)

0019151.0

7 AUG 2000

07AUG00 E558428-5 D00346

P01/7700 0.00-0019151.0

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Pace Micro Technology Plc

Victoria Road
Saltaire
Shipley
BD18 3LF

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

U.K.

6905293.001

4. Title of the invention

Deferred Internet Page Reformatting

5. Name of your agent (if you have one)

Bailey Walsh & Co.

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

5, York Place
Leeds
LS1 2SD

Patents ADP number (if you know it)

224001

6. If you are declaring priority from one or more earlier patent applications, give the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / years)

7. If this application is divided or otherwise derived from an earlier UK application, the earlier application

Number of earlier application

Date of filing
(day / month / years)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer "Yes" if:

Yes

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body
- See note (d)

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document.

Continuation sheets of this form

Description 6

Claim(s) 5

Abstract

Drawing(s)

10. If you are also filing any of the following, state how many of each item.

Priority Documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (Please specify)

11. I/We request the grant of a patent on the basis of this application

Signature

Date

G. Wood

04.08.00

12. Name and daytime telephone number of person to contact in the United Kingdom

G. Wood
0113 2433824

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help filling in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in this form you must remember to sign and date it.
- For details of the fee and ways to pay, please contact the Patent Office.

Deferred Internet Page Reformatting

The invention to which this application relates is an improved method for the generation of pages for an internet website on a display screen and particularly, although not necessarily exclusively, to the generation of pages using apparatus with relatively limited processing capability.

When fetching pages, web browsers tend to display the fetched information as soon as they can. On devices with processor power and random access memory (RAM) spare, this is typically the best approach, with whole-window backing stores in some browsers being used to avoid flicker. On devices such as broadcast data receivers (STBs), where processor power and RAM are both very limited and where a user may be trying to navigate through links on the page via a remote control handset rather than a mouse controller, frequent reformatting is much less desirable. Processor time is given up to the reformat process which can lead to jerky and unreliable navigation for the user, and there is often no method of backing store to alleviate flicker as the reformatted page is redrawn. Even on powerful desktop systems with backing stores, the frequent "reshuffling" effects seen when fetching some web sites can get annoying for the user.

Traditionally, web browsers reformat a web page whenever new page or image data arrives. Rapid reformats increase processor usage during page fetching, lead to more redraws and can hinder user navigation of the page via highlight-based (e.g. infra-red handset) systems. Instead of reformatting immediately, this invention describes an alternative system.

In a first aspect of the invention there is provided a deferring system which stops web browsers reformatting pages during a fetch of data each time the page layout needs to change, thus reducing processor overhead and the potential for flicker during fetching, such that the web browser is made to reformat at no more often than a certain fixed interval, rather than whenever new page or image data arrives.

This invention solves such problems through the reformat delay. When events happen that would normally cause an immediate reformat, the browser takes note of the highest point in the page that would be affected and starts a timer. Until the timer reaches zero or the entire page fetch is completed, which determination is based on the high water mark record, the parts of the page that have changed are reformatted. The number of reformat that typically occur over a page fetch, particularly when using a modem or if the HTML source of the page includes images but gives no indication of their size, can be reduced. This in turn means that, in total, less processor power is used during the fetch process and it proceeds more smoothly from the user's perspective.

A specific embodiment of the invention is now described as follows:

Internet web browsers usually can be required to reformat a page during fetching and generation for several reasons, two of which are:-

- 1) New page data arrives. This is particularly problematic for pages based heavily on HTML tables.

- 2) New image data arrives. If the HTML including the images gives no indication of their sizes, the browser has to use a standard sized placeholder and then reformat the page to fit the actual image size once enough of the image data is fetched to work its dimensions out.

This invention solves problems associated with these two reasons. Repeated reformatting, particularly if data is arriving in many small groups punctuated by short delays, can lead to a lot of flicker if the device has insufficient RAM and/or processor power to implement a whole-page backing store system. Many trivial reformats also tend to be more expensive to perform than one larger scope reformat encompassing all the changes that the individual smaller reformats were achieving.

In addition, if a user is navigating the page with a highlight-based device (e.g. not a mouse – cursor keys on a keyboard or IR handset are typical) the repeated reformats can make the navigation on images on the page displayed jerky. It can also become unreliable; for example, the user might be about to select a link when the part of the page it was in is reformatted and which can lead to an incorrect choice being selected. The deferred reformat model in accordance with the invention reduces the chance of the frequent formatting.

On modem based internet devices, small groups of page data followed by small delays are often seen. Case (1) above can then become particularly problematic, especially if the data arriving is part of a large table that takes large amounts of processor power to resize each time data is added. On devices with faster links, the page data size is often fairly insignificant compared to the link speed but this causes another problem, as in case (2) above. Images often fetch in rapid succession, and

if the HTML including those images gives no indication of their sizes then rapid, repeated reformatting can take place with traditional web browser layout methods.

In accordance with the invention, the reformatting is deferred. When a reformat is required for the first group of data fetched for a page, it is carried out immediately to give the user the earliest possible indication of progress. Next time the page needs to be reformatted for any reason, however, the y-coordinate of the topmost part of the page that is affected is stored and a timer is started. No actual reformat occurs. Once the time is running, any subsequent required reforms are also not carried out. They are instead deferred. The topmost affected y-coordinate is checked in all cases and the highest value recorded.

When the timer reaches zero, the browser carries out a page reformat from the highest affected point, where possible. Some web browsers are restricted by internal design to only be able to lay out all of the page data fetched so far in one go, and in these cases this is what will be done, but this is not as efficient as only reformatting those parts that have changed. All of the changes that were being deferred during the timer period will be included at this time. All new page data fetched so far gets included, avoiding multiple reforms as each chunk of new data comes in, and all image sizes known up to that point are incorporated in the new layout.

The value of the timer is chosen to balance frequency of reformatting against showing the user progress in the page fetch. The timer is preferably chosen on a per-application, not per-page basis; for example, a slow link may work better with a long timer, so as much page data as possible is assembled before a

reformat, whereas a fast link may work better with a short timer, but not so short that a flicker problem reappears. In use and as an example it is found that values of much less than four seconds tend to have little benefit over no delay at all, and values of more than fifteen to twenty seconds may make users believe the fetch has stalled. Since the value is "not" critical, the software providing the time service does not have to provide a particularly accurate time, which can be useful for operating systems where there is a limit on the number of claimants for high accuracy timers.

The timer is cancelled and an immediate reformat is performed if the page fetch completes with all images. It is possible that it would be desirable to also do this if the user activates an equivalent of a "stop" button for the fetch. Generally, however, it proves more confusing for the user to see a burst of activity just as they hit "stop" than it does for all activity to cease, even though this means that some of the data fetched for the page will not be displayed.

One example is the news site <http://cnn.com/> which, like many others of the larger, more busy sites, uses a large tables-based layout with the HTML including many images, some of which have no size specified for them. In use, without a deferred reformatting system or backing store there can be considerable flicker whilst the page fetches. Some browsers even wait until all the unknown size images have arrived, which leaves the user staring at very little for a long time. Sites such as this would benefit considerably from deferring reformat in accordance with the invention to, say, 10 second intervals at the most frequent, in accordance with the present invention.

Reformats occur less often and encompass more changes using a deferred model over a traditional model. Consequently, deferred reformatting can prove sufficiently beneficial that in embedded devices it becomes possible to select a lower power processor than would otherwise be required, reducing the cost of the hardware. The user of the device sees less flicker and smoother highlight-based navigation during page fetches, making the device more appealing to consumers.